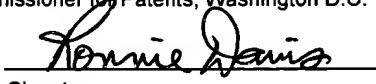APPLICATION FOR UNITED STATES LETTERS PATENT

for

Software-Implemented Method

for Identifying Nodes on a Network

by

Scott H. Hutchinson, Gregory M. Hanka

EXPRESS MAIL MAILING LABEL	
NUMBER	*EI 371 159 796 US*
DATE OF DEPOSIT	*January 20, 1999*

I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" service under 37 C.F.R. 1.10 on the date indicated above and is addressed to: Assistant Commissioner for Patents, Washington D.C. 20231.

_____
Signature

SOFTWARE-IMPLEMENTED METHOD
FOR IDENTIFYING NODES ON A NETWORK

1. BACKGROUND OF THE INVENTION

1.1. Field of the Invention

The invention was made in attempting to solve a specific problem in connection with "auditing" nodes, e.g., computer workstations and other computers (referred to sometimes here as microcomputers) on a computer network. The problem being addressed by the inventors was that of uniquely identifying nodes on a network for the purpose of maintaining a central database reflecting the hardware and software configurations of the respective nodes.

1.2. Description of Related Art

Recent years have witnessed the development of a category of software application products which enable a network system administrator ("SYSADMIN") to track computers and similar equipment ("nodes"), and their components, on computer networks. Applications in this category are sometimes referred to as "asset management" products.

Typically, asset management products assess the hardware and software components associated with a node on a network and maintain a central database of those nodes and components. The central database is usually remote (on a central computer) from the particular nodes being audited, inventoried or tracked; it is typically used by the SYSADMIN in managing network equipment and software.

An illustrative network on which such an asset management product might be used is shown in Figure 1. The network 100 includes two nodes 101, sometimes referred to as "client nodes," and a node 102, sometimes referred to as a "server node," connected by communications links 103.

1   The client nodes 101 and server node 102 are typically (but not necessarily) pro-

2   grammable computers. The two depicted client nodes 101 and the server node 102 are

3   merely representative examples; a typical network may include many such nodes.

4   The network 100 may be wide or local in geographical scope, *i.e.*, either a local

5   area network ("LAN") or a ("WAN"). Thus, the client nodes 101 and server node 102

6   may be geographically close, *e.g.*, in the same building, or geographically dispersed, *e.g.*,

7   in separate cities.

8   The network may employ any one of a variety of topologies, protocols, and ma-

9   chine architectures as are known in the art. For example, the network 100 may embody a

10   bus, star, or ring topology, may employ an Ethernet or token-ring protocol, or any other

11   type of network, and may employ a peer-to-peer or client-server architecture.

12   The communications links 103 may be optical fibers, electrical cables, telephone

13   wires, and/or wireless links depending on the particular embodiment implemented.

14   The foregoing examples are mentioned simply for purposes of illustration; those

15   of ordinary skill having the benefit of this disclosure will realize that the network 100

16   may take many other possible conventional forms.

17   The various client nodes 101 typically will either be programmable computers

18   (e.g., a user workstation) or will include one or more programmable processors (e.g., a

19   printer server). As such, each client node 101 will normally include writeable stor-

20   age 104, which may take the form of some or all of, e.g., a floppy disk drive, a hard disk

21   drive, removable storage media (e.g., a ZIP™ drive, JAZ™ drive, a writeable CD-ROM

22   drive, etc.), a tape drive, a flash-memory storage device, or any other suitable storage

23   medium now known or later developed.

24   Like the client node 101, the server node 102 contains a storage 105, e.g., a disk

25   drive.

26   Each client node 101 and the server node 102 contains a network interface card

27   (NIC) 107.

28

1.3.   The Desirability of Unique Node Identification

One task of an asset management product is to identify nodes uniquely and to recognize both when nodes 101 have been identified before and when they have not been, so as to recognize the node 101 each time the asset management product 'sees' it in the future, e.g., when the asset management product "audits" the node. This is required in order to match every node 101 up with its records in the central database. This allows the asset management product to know if there have been any changes in the components of a node 101 (*e.g.*, a floppy drive has been removed) since a previous audit.

2. SUMMARY OF THE INVENTION

The invention relates to an asset-management software product, which in one embodiment comprises a server program executing at a server node 102 and a client program executing at a client node 101.

In a first aspect of the invention, at the beginning of each audit, one or more unique attribute values (described in more detail below) of a client node 101 are detected by the client program. These attribute values are transmitted to the server program, which uses them to correlate the client node with a specific record in a central database (creating a new record if necessary). The one or more unique attribute values are also stored to a local database at the client node 101. Upon the next audit, the client program reads information from the local database to find out what the unique attribute values had been during the previous audit and transmits those values as well as the "new" detected values (which may now be different from the previous values). This transmission of out of date information as well as "new" information allows the server program to correctly correlate the client node with its (by now out of date) records in the central database if one or more of these unique attribute values of the node is changed.

In a second aspect of the invention, one specific attribute value tracked by the asset management product is the current address of the network interface card 107 ("NIC address") for each node 101 and any former NIC address it may have had in the past (i.e., its NIC address prior to obtaining the current one) for the purpose of node identification.

In a third aspect of the invention, one or more client nodes 101, referred to as "lonely nodes," either (1) has no active NIC 107 or (2) is configured so that the NIC address is undetectable by the client program. In the central database, the NIC address for each client node 101 is recorded for use during node identification; for lonely nodes, a fake NIC address is generated and stored. The fake NIC address is created in such a way that it is highly unlikely ever to duplicate any real NIC address in the network in question.

In a fourth aspect of the invention, the local database stored at a client node 101 is duplicated on multiple active partitions of its local hard-disk drive or drives 104, preferably on each such partition. Every copy of the local database receives a timestamp reflecting the time it was last updated, so that subsequent audits of the client node 101 can determine which copy (of possibly many) is the freshest.

3. BRIEF DESCRIPTION OF THE DRAWINGS

Other aspects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the drawings in which:

Figure 1 is a block diagram of a hypothetical prior-art network.

Figure 2 is a block diagram of a possible variation on such a network.

Figure 3 is a before-and-after block diagram of a node identification record in accordance with one implementation of the invention.

Figure 4 is a flow chart illustrating some operations performed in accordance with the first and second aspects of the invention described above.

Figure 5 shows a hypothetical three-way partitioning of a hard disk drive and redundant storage of a node identification record on each partition in accordance with the invention.

4. DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

4.1 Overview of the Problem

Essential as it is, unique identification in a network 100 is problematic for at least two reasons. First, computers in a typical network 100 are, from the point of view of a

1 software application, fundamentally alike, with only small differences setting them apart

2 from each other. Second, those small differences are subject to change over the course of

3 time and maintenance. There does not yet exist a standard, ubiquitous 'fingerprint' for

4 computers / nodes 101 yet, so asset management products must approximate one using

5 whatever shifting data they can find on each node.

6

4.1.1 The OEM Motherboard Serial Number as Potential Node Identifier

8 Referring to Figure 1, at this writing, the closest thing to a fingerprint for micro-

9 computers, and thus for a client node 101, is the new, fledgling standard for motherboard

10 serial numbers. A motherboard 108 is the comparatively large and complex circuit-board

11 on which most of a microcomputer's electronics are mounted. Of the many components

12 that make up a complete microcomputer, the motherboard 108 is the one best designated

13 the "central component," in the sense that system administrators expect to be able to

14 change any other component of a microcomputer (such as a hard-drive) and still have

15 their asset management product recognize that it is the same node as before (only now

16 with a new hard-drive). In the current psychology of the art, the motherboard 108 "is"

17 the microcomputer, all else is merely its interface to the world. So, node identification

18 would best track the motherboard 108, via some kind of unique serial-number built into it

19 by its manufacturer.

20 Unfortunately, as noted above, the standard for motherboard serial-numbers is a

21 fledgling standard. The standard is referred to as "Desktop Management Interface " or

22 "DMI." New computers do not necessarily support DMI completely, or support it at all,

23 to say nothing of the world of older computers built before DMI was finalized. So,

24 motherboard serial numbers are not likely for many years to suffice as a universal node

25 identification medium for an asset management product.

26

4.1.2 The NIC Address as Potential Node Identifier

28 Again referring to Figure 1, another unique node component for conventional mi-

29 crocomputer-based nodes is the network interface card 107 ("NIC"). As is well known to

30 those of ordinary skill, the term "network interface card" is a generic term for hardware

1   circuitry (usually with appropriate "firmware" programming) that provides an interface

2   between, e.g., a microcomputer and a network.   A network interface card 107 might take

3   the form of a circuit board that is inserted into an open slot inside a desktop computer or

4   in a docking station for a notebook computer; or a credit card-sized PCMCIA card (also

5   known as a PC Card) that can be inserted into a slot in a notebook computer; or circuitry

6   built directly into the motherboard of either a desktop or notebook computer, i.e., not as a

7   separate component at all.

8   Each NIC is assigned a unique built-in identifier by its manufacturer, known as

9   the media access control ("MAC") address, referred to hereafter as the "NIC address."

10   The NIC address is roughly analogous to the vehicle identification number ("VIN") as-

11   signed to cars and trucks by their manufacturers. Generally speaking, in this context an

12   "address" is a sequence of letters, numbers, etc., that serves as an identifier.

13   Referring to Figure 1, suppose that the server node 102 sends out data specifying

14   that the data is intended for the client node 101 having a NIC 107 whose NIC address is

15   e.g., "ABC123."   Normally, the data will be broadcast on the network 100. All client

16   nodes 101 will "hear" the data, including the client node 101 having a NIC 107 whose

17   NIC address is "QRS789," but only the specific client node 101 whose address is

18   "ABC123" will actually respond to the data.

19   The current framework within which NIC addresses are created is administered by

20   a committee of the Institute of Electrical and Electronics Engineers ("IEEE").   The

21   committee assigns blocks of addresses to each NIC manufacturer.   Each manufacturer

22   then permanently or semi-permanently installs specific addresses from its respective ad-

23   dress block(s) into one or more semiconductor chips on the NICs 107 by a process some-

24   times known as "burning" the address.

25   Ideally, the use of NIC addresses in this manner ensures that every NIC 107 in the

26   world has a unique address. (In this specification, the term "unique" is used in a local

27   sense, i.e., unique to a given network. NIC manufacturers sometimes reuse NIC ad-

28   dresses, but the odds are slim that two identically-addressed NICs 107 would ever show

29   up in the same network.)   Thus, since every computer on a conventional network has a

30   NIC 107, an asset management product could use the NIC address as its universal node

1   identification, and therefore each computer having a NIC 107 would have a unique ad-

2   dress on the network.

3   The NIC 107, however, is often not a permanent part of a microcomputer's moth-

4   erboard 108; very often it is a removable component plugged into the motherboard.

5   Changing a node's NIC 107 is sometimes required, if for example if the NIC becomes

6   defective, or if the network topology changes so as to require a different type of NIC, or if

7   the node itself is moved (*e.g.*, if a notebook computer is moved from docking station to

8   docking station, as shown in Figure 2, where each docking station contains its own sepa-

9   rate NIC). With a new NIC 107 presumably comes a new NIC address, and thus any as-

10 set management product relying solely on the NIC address for node identification will

11 falter when a node's NIC 107 changes in this way. (By analogy, the FBI would have a

12 similar problem if a person's fingerprints were to change every time the person got a

13 manicure.)

14

15 **4.1.3 The Hard Drive Contents as Potential Node Identifier**

16   Still another potentially unique component in a microcomputer is its fixed disk

17 drive ("hard disk" or "hard drive"), or more precisely, the contents thereof, shown as hard

18 disk storage 104 in Figure 1. During the first inspection or audit of a node, an asset man-

19 agement program can write its tracking data to a hidden file on a node's hard drive 104.

20 During subsequent inspections, the asset management program can retrieve the hidden

21 file and thus recognize the node as the one inspected earlier.

22   The potential downfalls of the hard-drive-contents approach are many, however.

23 For example:

24   1. Hard drives 104 are sometimes "reformatted," in which their entire contents

25 are erased to begin anew; in the process, any hidden files previously placed there by an

26 asset management program are normally lost.

27   2. Hard drives 104 are sometimes moved from one microcomputer (i.e., one cli-

28 ent node 101) to another, which can thoroughly confuse any asset management product

29 that presumes every hard drive 104 to be "married" to its motherboard 108.

-8-

1 3. It is not unusual for the contents of one hard drive 104 to be transferred in their

2 entirety to the hard drive 104 of another microcomputer, i.e., another client node 101; at

3 that point, two client nodes 101 have identical copies of the hidden file placed on the first

4 microcomputer / node by the asset management product, which will now incorrectly treat

5 both microcomputers / nodes 101 as the same node.

6 4. Nodes 101 with multiple operating systems installed will often have several

7 different file systems on their respective hard drives 104; thus, an asset management

8 product would "see" a different hard drive 104, in the sense of seeing a different file sys-

9 tem, each time the microcomputer / node 101 boots under a different operating system,

10 causing the asset management product to mistakenly regard each as a separate node.

11 * * *

12 There are other components of microcomputers that could be used as identifiers

13 for nodes 101, but unlike the three listed above, none of them are universally available

14 under all operating systems, or are available from all hardware manufacturers, or are ac-

15 ceptably likely to be unique within a given network.

16

17 **4.1.4 Summary of Selected Difficulties With Various Node-Identifier Ap-**

18 **proaches**

19 Some of the difficulties with various methods of identifying client nodes, includ-

20 ing prior-art approaches as well as aspects of the approaches disclosed and claimed in this

21 specification, are summarized in Table 1 below (in which X represents a failure case and

22 a forward slash / represents a nonfailure case):

NODE ID METHOD ↓	← FAILURE CASES →									
	Reformat C:	FDISK drive 0	New NIC	NIC swap	Multi-boot	New HDD	HDD swap	Diskless workstation	Cookie-cutter machine	GHOST™ machine
NIC Address	/	/	X	X	/	/	/	/	/	/
OEM serial No.	/	/	/	/	/	/	/	/	/	/
C: serial No.	X	X	/	/	X	X	X	X	X	X
Drive 0 boot record ID	/	X	/	/	/	X	X	X	/	X
Drive 0 firm-ware serial	/	/	/	/	/	X	X	X	/	/
Hidden file on boot drive	X	X	/	/	X	X	X	X	/	/
Hidden file on all drives	X	X	/	/	/	X	X	X	/	/

Table 1: Summary of Some Difficulties with Node-Identifier Approaches

NODE-IDENTIFICATION METHODS: The node-identification methods listed in the far-left column of Table 1 are as follows:

NIC address – as noted above, this is the "MAC address" burned into the firm-ware of the network interface card 107. It consists of six bytes, three for a vendor code assigned by IEEE and three for a serial number for use by that vendor. Vendors endeavor to avoid duplicating MAC addresses in their production NICs, sometimes even requesting additional vendor IDs from IEEE. In any given installation (network), it is safe to assume that all NIC addresses are unique.

OEM serial number – this is the serial number burned into the motherboard 108 of the workstation / node 101 by its manufacturer. With some difficulty, it can some-times be changed by a system administrator.

C: serial number – this is the four-byte serial number assigned to a formatted partition on a hard drive 104. It is recreated when the partition is reformatted, but other-wise does not change.

1 Drive 0 boot-record ID – this is a (for all intents and purposes) random number

2 created in the boot record of the "primary" hard drive 104 on the node's hard disk con-

3 troller, as indicated in the BIOS (basic input-output software) of the node 101. In many

4 computers, the drive 0 boot-record ID is created by an FDISK utility program at the time

5 that the "partitions" for the hard drive 104 are set up.

6 Drive 0 firmware serial – this is a serial number permanently burned into the ac-

7 tual hard drive unit 104, and in most cases, it is a very long, very unique string of charac-

8 ters assigned by the hard drive's manufacturer. Unfortunately, a few manufacturers do

9 not bother to use unique serial numbers.

10 Hidden file on boot drive – this is the practice of leaving a hidden file, e.g., an

11 .INI-type file, on the boot drive 104 of the workstation / node 101, containing node iden-

12 tification information.

13 Hidden file on all drives – this is the practice of leaving a hidden file, e.g., an

14 .INI-type file, on every available hard drive 104 on the workstation / node 101 in accor-

15 dance with the invention, as discussed below.

16

17 FAILURE CASES: The failure cases listed in the top row of Table 1 are the fol-

18 lowing:

19 Reformat C: – the primary hard drive partition of the workstation / node 101 was

20 reformatted, destroying the 'C: serial number' and also any hidden files (e.g., .INI-type

21 files) contained thereon.

22 FDISK drive 0 – a stronger version of 'Reformat C:' in which the workstation's

23 primary hard drive 104 was repartitioned and reformatted. Not only does this destroy the

24 'C: serial number' and any hidden files, e.g., .INI-type files, it can also reset the 'Drive 0

25 boot-record ID'.

26 New NIC – the workstation received a new network interface card 107, which

27 gives it a new NIC address.

28 NIC swap – the workstation / node 101 traded network interface cards 107 with

29 another workstation. Afterward, each workstation / node 101 has the other's former NIC

30 address.

1   Multi-boot – the workstation / node 101 uses boot-manager software (like System

2   Commander™) to boot different operating systems. The various operating systems may

3   designate different partitions as being the "C: drive," and even more commonly, may re-

4   gard different partitions as their boot drive. A further complication is that certain parti-

5   tions may be invisible or inaccessible on a particular operating system; for example, a

6   Windows NT™ NTFS™ partition cannot be accessed by DOS, Windows 95™, Win-

7   dows 98, Windows 2000, OS/2™.

8   New HDD – the workstation / node 101 received an entirely new hard drive unit

9   104. This condition is equivalent to 'FDISK drive 0', with the added complication that

10   the 'Drive 0 firmware serial' also changes.

11   HDD swap – the workstation traded hard drives with another workstation. After-

12   ward, each workstation has the other's former 'C: serial number', 'Drive 0 partition ID',

13   'Drive 0 firmware serial', and all hidden files, e.g., .INI-type files.

14   Diskless workstation – the workstation has no local hard drives, and hence, no 'C:

15   serial number', no 'Drive 0 partition ID', no 'Drive 0 firmware serial', and no possibility

16   of any hidden files, e.g., .INI-type files.

17   Cookie-cutter machines – the workstation was created from a prerecorded image

18   of a hard-drive, including an entire operating system and support software. As a result, it

19   has the same 'C: serial number' as all of its siblings.

20   GHOST™ machines – the workstation / node 101 was created using a PC imag-

21   ing program such as GHOST™, DiskImage, DiskIone, or other automatic software instal-

22   lation programs, which very thoroughly transplant the contents of one workstation's hard

23   drive 104 to the hard drive 104 of another workstation / node 101. As a result, the new

24   workstation / node 101 has the same "C: serial number" and the same "Drive 0 boot-

25   record ID" as all of its "siblings" created in this way.

26

27   **4.1.5 Inferences from Analysis**

28   From the information in Table 1, the following may be inferred:

"OEM serial number," i.e., a unique identification number of the motherboard 108, is the only 100% reliable node-ID method. Unfortunately, OEM serial-number detection is not yet widely available, and is far from an industry standard.

"NIC address" is an excellent alternative node-ID method, if only movements of NICs 107 could be handled somehow.

If "NIC address" could somehow be combined with "hidden file [e.g., .INI-type file] on all drives", the result would be 100% reliable for any single failure condition. (The case of multiple coincident failures is likely to be too complex to handle with anything other than the "OEM serial" approach.)

4.2. Illustrative Software-Based Solution

The multi-faceted approach of the invention is explained with reference to the network 100 shown in Figure 1.

4.2.1 Client Program; Server Software

The software running in each client node 101 includes an "agent" program referred to sometimes as an "audit" program and referred to here as a client program. The client program may be designed to run as a conventional foreground program, or as a background application, in whatever form is appropriate for the operating system in question (e.g., a terminate-and-stay-resident [TSR] program under MS-DOS or PC DOS, or a background service under other operating systems). Some well-known operating systems at this writing include, e.g., Windows 3.1; Windows 95; Windows 98; Windows NT; Windows 2000; Mac OS; various flavors of UNIX; and the like.

The client program exchanges information, via the communications links 103, with a server program that is likewise running on the server node 102. The server program performs many of the functions described below.

(In this specification, phrases such as "the client program doing X," where X is some function or functions, will be understood by those of ordinary skill as referring to one or more programmable processors performing the specified function(s) under control of the software in question.)

1   The client program can also run on the server node 102, so that the server program

2   can keep track of the hardware comprising the server node 102 itself. In that sense, the

3   server node 102 is also a client node 101.

4

4.2.2 Node-Identification Record at Client Node

6   Referring to Figure 3, the client program running on each client node 101 main-

7   tains a node-identification record 305 in a local database at the storage 104 of the node.

8   The node-identification record 305 may be stored in a separate file, e.g., an .INI-type file,

9   in the storage 104, or it may be added to an existing file, in either case as text information

10   as illustrated in the hypothetical example shown in Figure 3. (Still another alternative is

11   to store the node-identification record 305 in the Windows registry.) The file preferably

12   has appropriate attributes set in the usual manner so that the file is "hidden" from users.

13   It will be apparent that the selection of an .INI-type file in Figure 3 is for convenience

14   only and that other types of local data storage (again preferably hidden) may be used.

15   In one specific embodiment, each node-identification record 305 is stored in its

16   own .INI-type file having a unique fully-qualified file path (i.e., the complete "name" of

17   the file) and a timestamp indicating the date and time at which the file was last modified

18   (both the "name" and the timestamp are conventionally provided by the operating sys-

19   tem). The local database thus consists of whatever .INI-type files of that kind have been

20   created in this manner.

21   The node-identification record 305 preferably includes, possibly among other in-

22   formation, the value of one or more node-identification attributes of the node, i.e., at-

23   tributes of the hardware and/or software configuration of the node that tend to be unique

24   within a given network. For example, the node-identification record 305 may include

25   (i) a "previously-detected" NIC address, i.e., the NIC address detected by the client pro-

26   gram during the immediately-preceding audit (sometimes referred to in the appendixes as

27   the current NIC address stored in an .INI file), or, if no such address was detected, a

28   "fake" NIC address as described below; and (ii) a "former" NIC address, i.e., the most re-

29   cent NIC address detected by the client program that is different from the "previously-

30   detected" NIC address. The previously-detected NIC address is used for back-up pur-

1   poses in case no NIC is detected by the client program.   It will be appreciated that any

2   number of former NIC addresses may be stored in the node-identification record 305 if

3   desired, thus creating a history for that particular client node 101.

4   The node-identification record 305 may be initialized by the client program when

5   that software runs on the client node 101 for the first time. It may be updated either on a

6   scheduled basis or in response to specific events (e.g., every time the client program is

7   "booted up," i.e., started, or every time the client program performs an audit of the client

8   node 101).

9

10   ### 4.2.3 Central Database at Server Node

11   A central database (not specifically shown in the drawing figures) is stored in the

12   storage 105 at the server node 102.   Generally speaking, the database is a compilation or

13   some or all of the information stored in the node-identification records 305 at the respec-

14   tive client nodes 101, typically with one data record in the database per client node 101.

15   The database may also store other audit-related information provided by a client node

16   101, again typically in one record per node. The database is conventionally initialized

17   and periodically updated, either on a regular scheduled basis or in response to specific

18   events.

19

20   ### 4.3 Basic Node Identification Method

21   A node's NIC address represents a reliable client node 101 identification method.

22   The caveat, however, is that network interface card movements must be tracked some-

23   how. The node-identification records 305 and the central database provide tools that can

24   be used in such auditing.

25

26   ### 4.3.1 Initial Audit

27   During an initial audit of a client node 101, the client program running on that cli-

28   ent node conventionally detects the node's NIC address. The just-detected NIC address is

29   then stored in a new node-identification record 305 (e.g., a new, timestamped .INI-type

30   file) at the client node's data storage 104 as described in Section 4.2.2 above. If a NIC

1   address is not successfully detected, then a "fake" NIC address is used and stored instead,

2   as discussed in more detail in Section 4.5 below. The value of the just-detected NIC ad-

3   dress (or of the the "fake" NIC address) is written to the node-identification record 305

4   both as the "previously-detected" NIC address and as the "former" NIC address.

5

4.3.2 Transmission of Initial Audit Information

7   The client program then transmits, to the server program via the network 100, the

8   desired node-identification information for that client node 101, including two specific

9   items: The just-detected NIC address (or alternatively the "fake" NIC address), plus the

10 "former" NIC address from the new node-identification record 305. (Since this is the

11 first time the node has been audited, the "former" NIC address field will be empty; it may

12 be transmitted as a prearranged empty-field value, e.g., all zeros, or alternatively it may

13 be transmitted as a signal indicating the absence of a former NIC address.) The transmit-

14 ted node-identification information may also include, e.g., the OEM motherboard serial

15 number as discussed in Section 4.1, to the extent available.

16 When received by the server program, the fingerprint information is stored in the

17 database (e.g., by the server program itself or by a separate database management system

18 [DBMS] routine in response to a call from the server program).

19

4.3.3 Subsequent Audits

21 Referring to Figure 4: During subsequent audits of the same client node 101, the

22 client program again attempts to detect a NIC address (block 405). In addition, the client

23 program reads the most recent node-identification record 305, e.g., the .INI-type file with

24 the most recent timestamp.

25 If a NIC address is successfully detected, the address is compared to the contents

26 of the node-identification record 305 (block 410). If the node-identification record 305

27 contains a previously-detected node address that is identical to the just-detected node ad-

28 dress, then the client program knows that the NIC 107 has not changed since the last

29 audit, and therefore the node-identification record 305 is current and may be transmitted

30 with the former NIC address (block 415). However, if the "previously-detected" NIC ad-

1   dress in the node-identification record is different from the just-detected NIC address,

2   then the client program knows that the NIC 107 has changed since the last audit. The cli-

3   ent program therefore moves the "previously-detected" NIC address in the node-

4   identification record to the "former" NIC address field, and then copies the just-detected

5   NIC address into the "previously-detected" NIC address field in an updated version of the

6   node-identification record, identified with reference numeral 310 (block 420); the de-

7   tected NIC address and the newly-updated former NIC address are then transmitted

8   (block 425).

9   If, on the other hand, no NIC address was successfully detected, the client pro-

10 gram uses the "previously detected" NIC address from the node-identification record 305

11 as the "just detected" NIC address, assuming that the NIC has not changed since it was

12 last detected (block 430). If no "previously detected" NIC address is available in any

13 node-identification record 305, the client program generates a fake NIC address instead as

14 described in more detail in Section 4.5.

15 Once again the client program transmits desired node-identification information to

16 the server program as described in Section 4.3.2 above. The server program uses this in-

17 formation to locate the node's record in the central database. The server program seeks

18 the record according to the just-detected NIC address (which may in fact be "fake" or

19 "previously detected") and the "former" NIC address. If no record is found matching the

20 two, the server program seeks the record according to the "former" NIC address, on sus-

21 picion that the node's NIC address has recently changed and is therefore still recorded in

22 the central database under its "former" NIC address.

23 In the hypothetical case shown in Figures 3 and 4, the client node's NIC address

24 has just changed from "ABC123" to "DEF789." The server program will look up the

25 node as "DEF789 formerly ABC123," which will fail, so the server will then re-try

26 looking up the node, this time as "ABC123", which will succeed since ABC123 was in-

27 deed the node's NIC address during its last audit.

28 The central database is then conventionally updated to reflect the most recently

29 detected NIC address (be it real or fake, the server program does not care) — in effect,

30 now identifying the client node 101 as "the client node DEF789 formerly ABC123".

1 A pseudocode appendix setting forth the server-program algorithm in detail is re-

2 produced as Appendix 1 below.

3

4 **4.4 Replicated Node-Identification Records**

5 It was noted above that a local database is maintained at the data store 104 of each

6 client node 101. The local database contains a description of some of that client node's

7 unique attribute values that may be used for identification.

8 Referring to Figure 5 as a hypothetical example: In one aspect of the invention,

9 when the data store 104 includes one or more hard disks or similar partitionable storage

10 media, a mirrored copy of the local database is maintained on each active partition 505,

11 510, 515, etc., of each such hard disk. At the beginning of an audit, the client program

12 checks the respective internal timestamps of all accessible copies of the local database

13 505, 510, 515, etc., to determine which copy is most recent. The most recent copy is

14 utilized during the audit as described above. Afterward, the updated local database is re-

15 written onto the accessible partitions 505, 510, 515, etc., of all hard drives 104, overwrit-

16 ing any old copies. Included in the rewrite is an update of the timestamps of the copies.

17 The aforementioned functionality permits the asset management product to deal

18 with nodes that "boot" with multiple operating systems. Sometimes, a computer will

19 have several operating systems installed, and will boot between them at-will. Local hard-

20 disk partitions 505, 510, 515, etc., that are visible under one operating system are some-

21 times invisible under other operating systems. So, an audit under one operating system

22 may generate six copies of the local database, but then a subsequent audit under another

23 operating system may update only three of them. The next audit under the original op-

24 erating system will show three old copies and three new copies. Thanks to the times-

25 tamps, the client program can tell which copy or copies contain the latest information

26 about the node's unique attribute values. This information is used by the central database

27 for node identification.

28 In the hypothetical example shown in Figure 5, the storage 104 at the client

29 node 101 is configured with a hard disk drive that is divided into separate partitions 505,

30 510, 515, etc., that may be logically configured as drives C, D, and E respectively, with

1 different, selectively-bootable operating systems on each partition. In the illustration of

2 Figure 5, the three logical drives are shown as being bootable into MS-DOS, UNIX, and

3 Windows NT respectively.

4 Such a partitioning can cause complications for the audit process. Suppose that

5 the client node 101 is "booted up" (i.e., started) into, say, MS-DOS on drive C, but the

6 node-identification record 305 is stored only on drive D containing, say, the UNIX op-

7 erating system. It will be quickly recognized by those of ordinary skill having the benefit

8 of this disclosure that, because of certain limitations in MS-DOS, any files stored on

9 drives D and E – which could include the node-identification record 305 – may be inac-

10 cessible to client program running under MS-DOS unless the CONFIG.SYS file for MS-

11 DOS is properly configured. (It will be apparent that client program appropriate for the

12 operating system actually booted must be provided on an accessible logical drive.)

13 In such a configuration, the node-identification record 305 may be replicated, i.e.,

14 redundantly stored within each of the logical drives C, D, and E, thus making its contents

15 available to client program no matter which operating system is booted. Consequently,

16 if the client node 101 is booted into MS-DOS from drive C, and drives D and E are inac-

17 cessible, the copy of the node-identification record 305 can still be updated by client pro-

18 gram running under MS-DOS; the updated copy is identified by reference numeral 310.

19 As noted above, however, that in turn presents another issue: If only the copy 310

20 of the node-identification record on drive C is updated (because by hypothesis drives D

21 and E are inaccessible), then the replicated copies of the node-identification record 305

22 on drives D and E may be out of date the next time that UNIX or Windows NT is booted.

23 This issue may be addressed by having the client program, upon its own booting, check

24 the timestamps in any accessible copies of the node-identification record 305 and 310 and

25 use only the latest one as described above.

26 (As used in this specification, including in the claims, the term "redundantly

27 stored" is not intended to be limited to the situation in which all instances, i.e., all copies

28 of the node-identification record, have been synchronized. That is, the term "redundant

29 storage" is intended to include, not exclude, the situation in which one copy of the record

30 has been updated but the other copies have not yet been updated, as shown in Figure 3.)

- 19 -

4.5 Tracking of "Lonely" Nodes

It is not always possible to detect a network interface card 107 on certain client nodes 101, thus precluding the use of a NIC address as the exclusive node identifier for that client node 101. For example, when a client node 101 that is part of a Windows NT™ network 100 is booted under DOS, it is likely that the network interface card 107 will not be active if no DOS drivers for the network interface card 107 are installed. Or, referring again to Figure 2, an undocked notebook computer 200 may have its network interface card 107 located in a docking station 201; the notebook computer 200 is considered part of the network 100, but it does not have an active network interface card 107 because it is undocked. Other client nodes 101 in the network might not have a network interface card 107 at all, but they are still "in the network" 100, so to speak, from the perspective of the system administrator, who desires to be able to keep track of them. Such client nodes 101 are referred to as "lonely nodes." Without a NIC address, these nodes do not have the universal 'fingerprint' required by the asset management product for the recognition purposes.

So, when initially building a node-identification record 305 for a "lonely" client node 101 (i.e., one without an active network interface card 107), the client program generates, and stores in the node-identification record 305, a "fake" NIC address to correspond only to that particular lonely node. The generated NIC address is fake in the sense that it is unique to the network 100.

The fake NIC address is used by the client program and the server program in the same general way as a real NIC address, unless and until a real NIC address is subsequently detected on the node in question. If a real NIC address is subsequently detected on the node, then the fake NIC address is retired as the node's "former NIC address."

Fake NIC addresses must not duplicate any possible real NIC address, lest a fake NIC address accidentally duplicate the address of a real NIC on the network and thus confuse the two nodes in the central database. Non-duplication is accomplished by using a block of NIC addresses allocated to the applicant by IEEE. The asset management prod-

H: 308075(6LPN02!.DOC)

1   uct can create fake NIC addresses anywhere within this block under IEEE's guarantee

2   that no NIC manufacturer can be assigning NIC addresses in the same range.

3   In one implementation, the fake NIC address is generated by combining a known

4   signature portion (*e.g.*, a three byte signature code or NIC vendor ID assigned to the

5   software vendor by the IEEE) with a pseudorandomly generated portion, for a total of six

6   bytes of data. A typical fake NIC address looks like this: 00-90-D4-1F-E3-22. The first

7   three bytes of the fake NIC address consist of a NIC vendor ID assigned by the IEEE.

8   The last three bytes are generated pseudorandomly by the asset management product. An

9   example of an algorithm for generating pseudorandom portions is shown in Appendix 3.

10   The signature portion of the fake NIC address is included so that the server soft-

11   ware running on the server 102 will recognize that the NIC address was artificially gen-

12   erated. This provides the system administrator with greater node inventory reliability be-

13   cause the asset management product knows not to report the NIC address as a genuine

14   one. The pseudorandom portion is added to the fake NIC address 24 in case the network

15   100 has more than one lonely client node 101.

16   It will be appreciated by those skilled in the art that a conventional pseudoran-

17   dom-number generator can be used to generate the pseudorandom portion of the fake NIC

18   address. The use of a pseudorandom-number generator provides a reasonable assurance

19   that the fake NIC address will be unique within the network 100.   One algorithm for gen-

20   erating fake NIC addresses is set out in Appendix 3 below. Because each byte has 256

21   possible values, this algorithm yields 256 x 256 x 256 = 16,777,216 different possible

22   fake NIC addresses, which means the chance of getting duplicate fake NIC addresses is

23   acceptably small.

24   An additional advantage provided by the use of fake NIC addresses is that it

25   permits other existing asset management applications software to continue operating

26   normally without the system administrator having to worry about whether a client

27   node 101 actually has a NIC address.

28   Upon subsequent audits of the network 100, assuming that the lonely node is

29   eventually connected to the network via a network interface card 107, the lonely node's

30   actual NIC address will eventually be detected and the respective records will be updated

H: 308075(6LPN021.DOC)

1 in the node-identification record 305 at the client node itself and in a record at the data-

2 base at the server node 102. As noted above in this Section, the fake NIC address is then

3 retired as the node's now-former NIC address.

4

4.6 Some Alternative Implementations

6 In hindsight, it will be appreciated by those of ordinary skill having the benefit of

7 this disclosure that the node identification technique disclosed here can be used in a vari-

8 ety of situations.

9 For example, the node identification technique can be used any time information

10 is to be transmitted from a client node 101 to one or more other nodes, whether or not in

11 response to a query by the other node(s). For example, in hindsight it will be apparent

12 that the client node 101 may be programmed automatically to send reports of various

13 kinds to another node running appropriate server software, without waiting for an audit

14 command or other query from the other node.

15 As another example, the basic approach in which the client program sends out

16 both current and former identification information to a server program may be used in

17 contexts not involving a NIC address, e.g., over the Internet.

18 As still another example, the client program may be designed to perform some of

19 the functions of the server program, thus possibly freeing up the server program and its

20 host computer from operations that can be performed at the client program. This permits

21 the asset management product to operate on a standalone basis to that extent, with the

22 relevant portions of the database 204 being maintained (or replicated) at the local stor-

23 age 104.

24

4.7 Program Storage Device

26 It will be apparent to those of ordinary skill having the benefit of this disclosure

27 that the client program and the server software may be implemented by programming one

28 or more suitable general-purpose computers having appropriate hardware. The pro-

29 gramming may be accomplished through the use of one or more program storage devices

30 readable by the computer and encoding one or more programs of instructions executable

H: 308075(6LPN02!.DOC)

1 by the computer for performing the operations described above. The program storage

2 device may take the form of, e.g., one or more floppy disks; a CD ROM or other optical

3 disk; a magnetic tape; a read-only memory chip (ROM); and other forms of the kind well-

4 known in the art or subsequently developed. The program of instructions may be "object

5 code," i.e., in binary form that is executable more-or-less directly by the computer; in

6 "source code" that requires compilation or interpretation before execution; or in some in-

7 termediate form such as partially compiled code. The precise forms of the program stor-

8 age device and of the encoding of instructions is immaterial here.

9

10 ## 5. ORDER OF OPERATIONS IN METHOD CLAIMS

11 Some of the claims below recite the performance of certain operations or func-

12 tions. It will be understood, by those of ordinary skill having the benefit of this disclo-

13 sure, that the operations or functions in question are not necessarily required to be per-

14 formed in the specific order in which they are listed in the claims.

15

16 ## 6. SOFTWARE PSEUDOCODE APPENDIXES

17 The appendixes below are pseudocode listings for a specific implementation of

18 the invention by the assignee in client program and server software.

H: 308075(6LPN021.DOC)